arXiv:1703.10363v1 [math.OC] 30 Mar 2017

Estimating effective connectivity in linear brain network models

G. Prando, M. Zorzi, A. Bertoldo and A. Chiuso †

***Abstract*— Contemporary neuroscience has embraced network science to study the complex and self-organized structure of the human brain; one of the main outstanding issues is that of inferring from measure data, chiefly functional Magnetic Resonance Imaging (fMRI), the so-called effective connectivity in brain networks, that is the existing interactions among neuronal populations. This inverse problem is complicated by the fact that the BOLD (Blood Oxygenation Level Dependent) signal measured by fMRI represent a dynamic and nonlinear transformation (the hemodynamic response) of neuronal activity. In this paper, we consider resting state (rs) fMRI data; building upon a linear population model of the BOLD signal and a stochastic linear DCM model, the model parameters are estimated through an EM-type iterative procedure, which alternately estimates the neuronal activity by means of the Rauch-Tung-Striebel (RTS) smoother, updates the connections among neuronal states and refines the parameters of the hemodynamic model; sparsity in the interconnection structure is favoured using an iteratively reweighting scheme.**

Experimental results using rs-fMRI data are shown demonstrating the effectiveness of our approach and comparison with state of the art routines (SPM12 toolbox) is provided.

I. Introduction

Recent research in the field of computational neuroscience has shown a growing interest in the estimation of the effective connectivity among brain regions, that is, the existing interactions between neuronal populations located in different brain areas [1]. What makes the problem particularly challenging is the lack of a direct measurement of the neuronal activity: standard neuroimaging techniques (such as EEG, MEG and fMRI) only provide an indirect measurement of such activity. In particular we shall be concerned with fMRI data, which provide noisy measurements of the so-called BOLD (Blood Oxygenation Level Dependent) signal. The dependency between this quantity and neuronal activity is described by the so-called hemodynamic response. In brief, an increment of the synaptic activity in a certain brain region causes an increased blood flow to that area, which in turn augments the ratio of oxygenated haemoglobin w.r.t. the deoxygenated one in that region. By exploiting the different magnetic properties of these two types of hemoglobin, fMRI is able to detect local changes in their relative concentration, thus highlighting the eventual presence of synaptic activations.

Existing methods tackle the problem of effective connectivity estimation by postulating a parametric generative model of the observed fMRI data, which not only accounts for the coupling among neuronal populations, but also for the aforementioned hemodynamic response. Nonlinear dynamical models are used and generally referred to as Dynamic Causal Models (DCMs) [2]: both deterministic [2], [3], [4] and stochastic [5], [6] formulations exist, with the latter being more suited to describe resting state fMRI data.

Having postulated a DCM, effective connectivity inference is reduced to model parameters estimation, which is typically accomplished using Variational Bayes techniques [7], [8], [9]. However, existing approaches are computationally intensive and are effective only on small brain networks. A large portion of the required computational effort is spent for the determination of the structure of the network connecting the different brain regions. Indeed, this choice involves the comparison of several candidate structures, making the problem combinatorial, [10]. In addition, the inversion of stochastic DCMs appears more involved since it also requires the estimation of the hidden random neuronal activity [9].

DCM is a Bayesian approach originally proposed to investigate from fMRI signal the influence on the neuronal system exerts by external stimuli, i.e. the so-called task-evoked effective connectivity. DCM can be also be used to study the directed connections among distributed brain areas when only endogenous brain fluctuations are present, i.e. the so-called resting-state effective connectivity. Our study will focus on this latter problem. We develop a linear DCM, and adapt filtering techniques to estimate the hidden neuronal activity as well as the DCM parameters. Differently from existing methodologies, which require the comparison of a set of candidate patterns of effective connectivity, our approach estimates directly the connectivity pattern, by resorting to an iteratively ℓ_1 reweighed algorithm which induces sparsity in the estimated connectivity matrix, thus automatically detecting the structure of the brain network.

The paper is organized as follows. We briefly review classical DCMs and parameters inference techniques in Sec. II, while in Sec. III we propose an approximated linear stochastic DCM, which is suited for resting state fMRI data. In Sec. IV we deals with the estimation of the DCM parameters (unrealistically) assuming to measure the neuronal activity. We deal with the realistic situation in which only the BOLD signal is available in Sec. V, where we illustrate the joint estimation of the DCM parameters and of the hidden neuronal activity. Using simulated fMRI data, we evaluate the effectiveness of the proposed method in Sec. VI, while we provide some concluding remarks in Sec. VII.

This work has been partially supported by the project "Statistical learning methods for estimating the effective connectivity of human brain networks" (BIRD162411/11).

† Dept. of Information Engineering, University of Padova (e-mail: {prandogi,zorzimat,bertoldo,chiuso}@dei.unipd.it)

II. Problem Formulation

A. Generative model

A DCM is a nonlinear multi input-multi output (MIMO) dynamic system, with one output for each monitored brain region. A DCM basically consists of two components: the first describes the dynamic coupling among neuronal populations, while the second maps the neuronal activity to the BOLD signal observed in each brain area, through the hemodynamic response. Accordingly, the system states are grouped in two classes: one contains the n neuronal states (one for each monitored brain region), while the other refers to the biophysical quantities involved in the hemodynamic model.

In the seminal paper [2], the coupling among the neuronal states $x \in \mathbb{R}^n$ is modelled through a deterministic bilinear interaction:

$$\dot{x}(t) = \left(A + \sum_{j=1}^{m} u_j(t) B_j \right) x(t) + Cu(t), \qquad t \in \mathbb{R} \quad (1)$$

where $u(t) = [u_1(t) \ \cdots \ u_m(t)]^T$, with $u_j(t)$ denoting the j-th designed external stimulus. Effective connectivity in the absence of external solicitations is encoded by the matrix $A \in \mathbb{R}^{n \times n}$, while the bilinear term $B_j \in \mathbb{R}^{n \times n}$ accounts for the change in the neuronal coupling due to the j-th input; finally, $C \in \mathbb{R}^{n \times m}$ modules the direct influence of external inputs on the neuronal activity. The system matrices $\theta_c = \{A, B_1, B_2, ..., B_m, C\}$ are unknown parameters which need to be estimated in order to retrieve the effective connectivity among the considered brain areas.

Extensions of the bilinear model (1) can be found in the literature, see e.g. [3] and [4].

In this paper we will consider the resting state activity, which can be modelled as [6]

$$\dot{x}(t) = Ax(t) + Cu(t) + w(t), \quad t \in \mathbb{R} \quad (2)$$

where $w(t)$ is a stochastic process which accounts for the unmeasurable random fluctuations that drive the resting state activity. Even if in this situation the exogenous inputs $u(t)$ are typically set to zero, model (2) can still accommodate the presence of external (non-modulatory) signals. In this paper we assume $u(t) = 0$.

The hemodynamic response is generally charachterized through an extension of the so-called Balloon-Windkessel model, [11], [12], a SISO dynamic system with the neural activity $x_i(t)$, i.e. the component of $x(t)$ in position i, as input and the corresponding observed BOLD signal $y_i(t)$ as output. The model is described by $4-th$ order nonlinear state space model depending upon the biophysical parameters $\theta_h = \{\kappa, \gamma, \tau, \alpha, \rho\}$ more specifically:

$$\dot{s}_i(t) = x_i(t) - \kappa_i s_i - \gamma_i (f_i(t) - 1), \qquad t \in \mathbb{R} \quad (3)$$

$$\dot{f}_i(t) = s_i(t) \quad (4)$$

$$\tau_i \dot{v}_i(t) = f_i(t) - v_i^{1/\alpha}(t) \quad (5)$$

$$\tau_i \dot{q}_i(t) = \frac{f_i}{\rho_i} \left(1 - (1-\rho_i)^{1/f_i(t)} \right) - v_i^{1/\alpha}(t) \frac{q_i(t)}{v_i(t)} \quad (6)$$

$$y_i(t) = V_0 \left(k_1 (1 - q_i(t)) + k_2 \left(1 - \frac{q_i(t)}{v_i(t)} \right) + k_3 (1 - v_i(t)) \right) + e_i(t) \quad (7)$$

The so called hemodynamic states $\{s_i, f_i, v_i, q_i\}$ are biophysical quantities which are affected by the neuronal activity: s_i denotes the vasodilatatory signal s_i, f_i is the blood inflow, v_i and q_i are respectively the blood volume and the deoxyhemoglobin content, $e_i(t)$ is the observation error, which may arise from thermal and physiological causes during data acquisition. The output equation (7) depends on the resting blood volume fraction V_0 (typically set to 0.4) and on the constants k_1, k_2 and k_3. Despite these have found different characterizations in the literature, a general form is derived in [13]:

$$k_1 = 4.3 \vartheta_0 \rho_i T_E, \qquad k_2 = \varepsilon r_0 \rho_i T_E, \qquad k_3 = 1 - \varepsilon. \quad (8)$$

Here, T_E is the echo time, while $\vartheta_0 = 40.3 s^{-1}$ denotes the frequency offset at the outher surface of the magnetized vessel for fully deoxygenated blood at 1.5 T; $\varepsilon \approx 0.4$ is the ratio of intra- and extravascualr signal, while $r_0 = 25 s^{-1}$ is the slope of the relation between the intravascular relaxation rate and oxygen saturation.

Collecting the neuronal and the hemodynamic states into the vector $z \in \mathbb{R}^{5n}$ and the two parameters sets in θ, i.e. $\theta = \{\theta_c, \theta_h\}$, a stochastic DCM can be compactly described as the following nonlinear dynamic system:

$$\begin{cases} \dot{z}(t) &= f(z(t), u(t), w(t), \theta) \\ y(t) &= g(z(t)) + e(t), \qquad t \in \mathbb{R} \end{cases} \quad (9)$$

where $y(t) = [y_1(t) \ \cdots \ y_n(t)]^T$ and similarly for $e(t)$. The practical implementation of (9) is as follows [1]. The simulated BOLD data, equivalently fMRI data, are given by discretizing the Balloon model (3)-(7) as well as the trajectory of the neural activity in (2) using the sampling time $T_R = 2$s, which is the time often used for the data acquisition. With some abuse of notation, $x(k) := x(kT_R)$ and $y(k) := y(kT_R)$, with $k \in \mathbb{Z}$, be, respectively, the discretized trajectory of x and y. By (2), with $u(t) = 0$, it is not difficult to see that

$$x(k+1) = e^{AT_R} x(k) + w(k) \quad (10)$$

where $w(k) \sim \mathcal{N}(0, Q)$ and

$$Q = \sigma^2 \int_0^{T_R} e^{A\tau} e^{A^T \tau} \mathrm{d}\tau. \quad (11)$$

B. Effective connectivity estimation

When considering the resting state condition, the effective connectivity estimation problem consists in estimating θ_c, i.e. matrix A (and possibly C), using the measured BOLD signal $y(k)$, $k = 1 \dots N$. It is worth noting that the estimation of θ_c requires to estimate θ_h as well, i.e. the parameters of the Balloon model. This appears as a challenging task, because of the nonlinearity of system (9), but also of the unavailability of the random states z, which have to be somehow inferred. Previous works have tackled this problem

[1] See `http://www.fil.ion.ucl.ac.uk/spm/`.

following two different routes.

A first approach, referred to as *spectral DCM*, [6], estimates the second-order statistics (i.e. cross-spectra) of the hidden states z, which are assumed to be stationary. In this way, the inference of the hidden time-series z is avoided. Furthermore, by parametrizing the spectral density of both process and measurement noise (typically as a power law distribution), the DCM is made deterministic and its parameters are estimated through standard Bayesian approaches [2], [8], postulating Gaussian priors on the parameters θ. Typically, a stability-inducing prior is formulated for the matrix A, while empirical priors are assigned to the hemodynamic parameters θ_h. However, the nonlinear nature of system (9) complicates the analytical computation of the posterior probability density function, as well as of the marginal likelihood. To overcome this bottleneck, a variational Bayes scheme is used, under the mean-field and the Laplace assumptions: namely, the approximated marginal posterior distributions are considered independent and Gaussian (see [7], [9] for details).

A second family of approaches still parametrizes the distribution of the process and observation noises in (9), but operates in the time-domain, thus estimating not only the model parameters θ but also the hidden neuronal states z. The most common methods are Dynamic Expectation Maximization (DEM) [8] and Generalized Filtering (GF) [14], [5]: even if both adopt the Variational Bayes procedure, DEM uses the mean-field and the Laplace approximations, while GF only exploits the latter.

Despite these techniques represent the state-of-the-art for the estimation of brain effective connectivity, computational feasibility limits their application to small networks (typically, no more than 10 regions). In particular, this drawback severely affects DEM and GF [15]. The reformulation of the nonlinear stochastic DCM (9) as a linear stochastic dynamic model may represent a way to address this issue. In this paper we pursue this idea, by deriving an approximated linear DCM according to the procedure detailed in Sec. III. This reformulation will allow us to jointly estimate the hidden neuronal states z and the DCM parameters θ. It is worth noting that a recent contribution in this direction focuses on deterministic DCMs and transforms the neuronal state equation (2) (with $w(t) = 0$) into the frequency domain; then, using a fixed hemodynamic response, the inversion of a deterministic DCM is formulated as a Bayesian linear regression problem, [16].

III. Statistical Linearization of DCMs

We model the hemodynamic response as a Finite Impulse Response (FIR) model which takes as input only a neuronal state $x_i(k)$ and as output the BOLD signal $y_i(k)$,

$$y_i(k) = \sum_{l=0}^{s-1} h_l x_i(k-l) \tag{12}$$

where the length s of the impulse response $h := [h_1 \ \cdots \ h_s]^T$ is chosen large enough to retain all the relevant dynamics components. A prior on the impulse response h is derived, using statistical linearisation techniques, from the Baloon model (3)-(7) where the parameters $\theta_h = \{\kappa, \gamma, \tau, \alpha, \rho\}$ are described by prior distributions typically used in the literature [2]. This will lead to linear model of the form $h = H\alpha$ (see eq. (14) below) where α will have to be estimated under a Bayesian perspective using the available fMRI data. To to so, we first compute a population of typical responses generated by the Balloon model (3)-(7), then we define h as the linear combination of their empirical mean and of the first p principal components of their sample covariance matrix. Specifically:

1) Sample $\theta_h^{(j)}$, $j = 1, ..., N_s$ from the empirical Gaussian distributions reported in Table 1 of the seminal paper [2].
2) For each $\theta_h^{(j)}$ compute the output of the Balloon model (3)-(7), where $x_i(k) = \delta(k)$ and $\delta(k)$ is the Kronecker delta function. Let $h^{(j)}$ be the corresponding output truncated at length s, which represents a statistic of h.
3) Compute the empirical mean $\bar{h} = \frac{1}{N_s} \sum_{j=1}^{N_s} h^{(j)}$.
4) Compute the empirical covariance matrix $\Sigma_h \in \mathbb{R}^{s \times s}$:

$$\Sigma_h = \frac{1}{N_s} \sum_{j=1}^{N_s} (h^{(j)} - \bar{h})(h^{(j)} - \bar{h})^T. \tag{13}$$

5) Compute the eigenvalues decomposition of Σ_h, $\Sigma_h = USU^T$, where $S := \mathrm{diag}(s_1, ..., s_s)$ and $U := [u_1 \ u_2 \ \cdots \ u_s]$.
6) Define h as

$$h := H\alpha = \begin{bmatrix} \bar{h} & u_1 & u_2 & \cdots & u_p \end{bmatrix} \alpha, \quad \alpha \in \mathbb{R}^{p+1}. \tag{14}$$

Figure 1 illustrates the empirical mean $\bar{h}$ of the sampled population of hemodynamic responses and the singular values of their sample covariance matrix Σ_h, showing that there is limited variability in the generated responses.

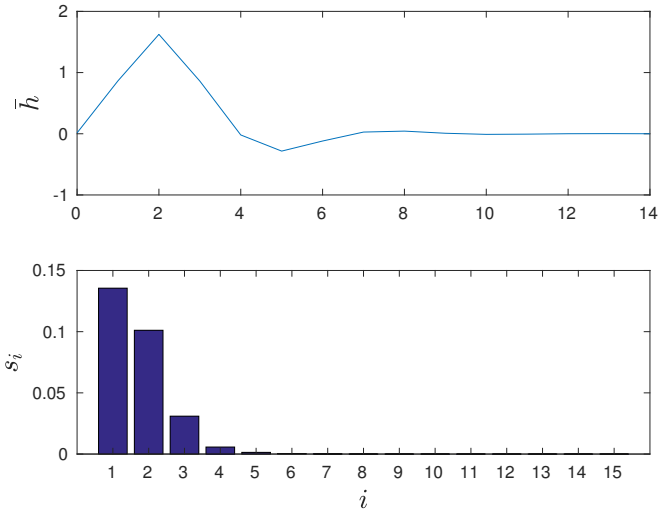

Fig. 1: *Top:* Empirical mean $\bar{h}$ of the sample hemodynamic responses. *Bottom:* Singular values s_i, $i = 1, ..., s$, of the sample covariance matrix Σ_h.

By introducing the extended state $\mathbf{x}(k) \in \mathbb{R}^{ns}$

$$\mathbf{x}(k) := \begin{bmatrix} x^T(k) & x^T(k-1) & \cdots & x^T(k-s+1) \end{bmatrix}^T, \quad (15)$$

the matrices

$$\mathbf{A} := \begin{bmatrix} e^{AT_R} & \mathbf{0} \\ I_{n(s-1)} & \mathbf{0} \end{bmatrix}, \qquad \mathbf{C} := (H\alpha)^T \otimes I_n, \quad (16)$$

and in view of (10), we can write the linear DCM as a classical stochastic state-space model

$$\begin{cases} \mathbf{x}(k+1) &= \mathbf{A}\mathbf{x}(k) + \mathbf{w}(k) \\ y(k) &= \mathbf{C}\mathbf{x}(k) + e(k). \end{cases} \quad (17)$$

Here, $\mathbf{w}(k) := [w^T(k)\ \mathbf{0}]^T \in \mathbb{R}^{ns}$ and we assume

$$\mathbf{w}(k) \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}), \qquad \mathbf{Q} := \text{blkdiag}(Q, \mathbf{0}) \quad (18)$$

$$e(k) \sim \mathcal{N}(\mathbf{0}, R), \qquad R := \lambda^2 I_n \quad (19)$$

where blkdiag$(\cdot)$ is the block-diagonal operator and the covariance matrix Q has been defined in (11). Notice that model (17) only depends on the neuronal states $x \in \mathbb{R}^n$ and on the parameters $\eta = \{A, \alpha, \sigma, \lambda\}$. Next sections will describe how these are estimated, starting from the simpler (but unrealistic) situation in which the neuronal states x are measurable, Sec. IV, to the situation in which x are not measurable, Sec. V.

IV. Parameters Estimation: Measured Neuronal Activity

In this Section we unrealistically assume that observations of the neuronal activity (encoded in the states x) are available and we detail how the parameters defining model (17) are estimated.

A. Estimation of matrix A

The matrix A is expected to be sparse, encoding the fact that each $x_i(t)$ is influenced by the activities of few other areas $x_j(t)$. To induce sparsity in the estimated A, we use a reweighted ℓ_1 method, [17], [18], originally proposed for linear regression models. However, because of the nonlinearity (in the parameters) of model (10), such approach cannot be straightforwardly applied but needs to be properly adapted. Introducing the matrices

$$X_+ := \begin{bmatrix} x^T(2) \\ \vdots \\ x^T(N) \end{bmatrix}, \ X := \begin{bmatrix} x^T(1) \\ \vdots \\ x^T(N-1) \end{bmatrix},$$

$$W := \begin{bmatrix} w^T(1) \\ \vdots \\ w^T(N-1) \end{bmatrix}. \quad (20)$$

we rewrite (10) in the non-linear regression form

$$X_+ = Xe^{A^T T_R} + W. \quad (21)$$

which can be linearised exploiting $e^{A^T T_R} \simeq I + A^T T_R$:

$$\Delta X = XA^T T_R + W \quad (22)$$

where $\Delta X = X_+ - X$. Using the vectorization operator, we can rewrite (22) as follows

$$\mathrm{x} = \Phi a + \mathrm{w} \quad (23)$$

where $\mathrm{x} := \text{vec}(\Delta X)$, $\Phi := (I \otimes X)T_R$, $a := \text{vec}(A^T)$, $\mathrm{w} := \text{vec}(W)$. Adopting the Sparse Bayesian Learning perspective [19], we model $a \sim \mathcal{N}(0, \Gamma)$ with $\Gamma := \text{diag}(\gamma_1 \dots \gamma_{n^2})$ and independent of w, so that the covariance matrix of x is $\text{Var}[\mathrm{x}] = \Phi\Gamma\Phi + Q \otimes I_N$. Direct application of the reweighted ℓ_1 method [17] to the approximated linear model (23) is not possible here, since $T_R = 2$s is far to be small, (23) is only a crude approximation of the continuous time model. For this reason, we propose Algorithm 1 where $a^{(j)} = \text{vec}(A^{(j)})$; $a_i^{(j)}$ is the entry in position i in vector $a^{(j)}$; $\Gamma^{(j)} = \text{diag}(\gamma_1^{(j)} \dots \gamma_{n^2}^{(j)})$; ϕ_i is the i-th column of Φ. Here the update for γ_i's is the same as [17, formula (29)] while the main difference rests in update for $A^{(j)}$, which is set as the posterior mean of (21), i.e. using the exact model, rather than the posterior mean of the approximated model (23); more precisely:

Algorithm 1 Proposed reweighted ℓ_1 method

Inputs: $x(k),\ k = 1, \dots, N$

Initialization: $A^{(0)} = -I_n$, $\gamma_i^{(0)} = 0.25\ i = 1 \dots n^2,\ j = 0$

1: **repeat**

2: $\quad \sigma^{(j)^2} = \frac{1}{Nn} \| X_+ - X e^{A^{(j)^T} T_R} \|^2_{Q^{-1}}$

3: $\quad Q = \sigma^{(j)^2} \int_0^{T_R} e^{A^{(j)}\tau} e^{A^{(j)^T}\tau} d\tau$

4: $\quad A^{(j+1)} = \underset{A \in \mathscr{A}}{\text{argmin}} \ \Big\{ \| X_+ - X e^{A^T T_R} \|^2_{Q^{-1}}$
$\qquad + \text{vec}(A^T)^T \left[\Gamma^{(j)} \right]^{-1} \text{vec}(A^T) \Big\}$

5: $\quad \gamma_i^{(j+1)} = (a_i^{(j+1)})^2 + \gamma_i^{(j)}$
$\qquad - (\gamma_i^{(j)})^2 \phi_i^T (\Phi\Gamma^{(j)}\Phi^T + Q \otimes I_N)^{-1} \phi_i,$
$\quad i = 1, \dots, n^2$

6: $\quad k = k + 1$

7: **until** $\|A^{(j)} - A^{(j-1)}\|_F / \|A^{(j)}\|_F$ is sufficiently small

Outputs: $A^{(j)}, \sigma^{(j)}$

$$a^{(j+1)} = \left(\Phi^T Q^{-1} \Phi + [\Gamma^{(j)}]^{-1} \right)^{-1} \Phi^T Q^{-1} \mathrm{x}. \quad (24)$$

Notice that, the search of the optimal $A^{(j)}$ is restricted to the set $\mathscr{A}$. In this paper, $\mathscr{A}$ is chosen as the set of stable matrices, i.e. with eigenvalues having negative real part.

Remark 1: It is worth to stress that the proposed reweighted scheme automatically selects a sparsity pattern for the matrix A. This represents a significant computational savings w.r.t. the Variational Bayes methods mentioned in Sec. II-B, which instead have to be provided with several candidate sparsity patterns, among which the best is selected resorting to model selection techniques, e.g. using the free-energy [10]; this gives rise to a combinatorial explosion of candidates making practical inference infeasible.

B. Estimation of the coefficients α

In Sec. III, a linear model for the hemodynamic response has been derived. As previously remarked, this depends on some parameters α that have to be estimated from data. Here we assume to have measured N data-pairs $(x(k), y(k))$ related by the output equation of model (17), that is

$$y(k) = \mathbf{C}\mathbf{x}(k) + e(k), \qquad k = 1,...,N. \tag{25}$$

The previous equation can be rewritten as

$$y(k) = \varphi(k)H\alpha + e(k), \qquad k = 1,...,N \tag{26}$$

by defining $\text{vec}(\varphi(k)) := \mathbf{x}(k)$ and observing that $\mathbf{C}\mathbf{x}(k) = \text{vec}(\varphi(k)H\alpha) = \varphi(k)H\alpha$. Consequently, introducing the matrices

$$Y := \begin{bmatrix} y(1) \\ \vdots \\ y(N) \end{bmatrix}, \quad \mathbf{H} := \begin{bmatrix} \varphi(1) \\ \vdots \\ \varphi(N) \end{bmatrix} H, \quad E := \begin{bmatrix} e(1) \\ \vdots \\ e(N) \end{bmatrix} \tag{27}$$

we have the following linear regression form

$$Y = \mathbf{H}\alpha + E. \tag{28}$$

We consider the Gaussian prior $\alpha \sim \mathcal{N}(\mu_\alpha, \Sigma_\alpha)$, where

$$\mu_\alpha := [1 \ 0 \ \cdots \ 0]^T, \qquad \Sigma_\alpha := \text{diag}(s_1,...,s_p) \tag{29}$$

and s_i, $i = 1,...,p$, are the first p singular values of the empirical covariance matrix Σ_h in equation (13). Recalling that $e(k) \sim \mathcal{N}(0, \lambda^2 I_n)$ and exploiting standard rules on conditioned Gaussian variables, we easily derive the MAP estimate of α as

$$\hat{\alpha} = \mu_\alpha + (\mathbf{H}^T\mathbf{H} + \hat{\lambda}^2\Sigma_\alpha^{-1})^{-1}\mathbf{H}^T(Y - \mathbf{H}\mu_\alpha) \tag{30}$$

where

$$\hat{\lambda}^2 = \frac{(Y - \mathbf{H}\hat{\alpha}_{LS})^T(Y - \mathbf{H}\hat{\alpha}_{LS})}{Nn - p - 1} \tag{31}$$

and $\hat{\alpha}_{LS} = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T Y$.

V. Parameters Estimation: Unmeasured Neuronal Activity

This section deals with the realistic situation in which only the resting-state fMRI data are observed and the neuronal states x are hidden. Sec. V-A will describe how we estimate x starting from the measured BOLD signal y, having fixed the model parameters $\eta = \{A, \alpha, \sigma, \lambda\}$; the subsequent Sec. V-B proposes two algorithms for the combined estimation of x and η.

A. Neuronal Activity Estimation

We assume that we are provided with N observations of the BOLD signal y, modelled according to the linear dynamic model (17), specified by the parameters η (here supposed to be known). We use the Rauch-Tung-Striebel smoother (RTSS) [20] to estimate the hidden neuronal states $x(k)$, $k = 1,...,N$, and in particular, the closed form smoothing solution to the linear model (17), that is

$$p(\mathbf{x}(k)|Y, \ \eta) = \mathcal{N}(\hat{\mathbf{x}}^s(k), \mathbf{P}^s(k)). \tag{32}$$

The implementation of the RTSS is summarized in Algorithm 2.

Algorithm 2 RTS Smoother

Inputs: $y(k)$, $k = 1,..,N$; $\mathbf{A}, \mathbf{C}$ in Eq. (16), $\mathbf{Q}$ and R in Eqs. (18), (19).

Forward Recursion

Initialize: $\hat{\mathbf{x}}(0) = 0$, $\mathbf{P}(0) = I_{ns}$
for $k = 1,...,N$ **do**
  $\hat{\mathbf{x}}^-(k) = \mathbf{A}\hat{\mathbf{x}}(k-1)$
  $\mathbf{P}^-(k) = \mathbf{A}\mathbf{P}(k-1)\mathbf{A}^T + \mathbf{Q}$
  $S(k) = \mathbf{C}\mathbf{P}^-(k)\mathbf{C} + R$
  $K(k) = \mathbf{P}^-(k)\mathbf{C}^T S^{-1}(k)$
  $\hat{\mathbf{x}}(k) = \hat{\mathbf{x}}^-(k) + K(k)[y(k) - \mathbf{C}\hat{\mathbf{x}}^-(k)]$
  $\mathbf{P}(k) = \mathbf{P}^-(k) - K(k)S(k)K^T(k)$

Backward recursion

Initialize: $\hat{\mathbf{x}}^s(N) = \hat{\mathbf{x}}(N)$, $\mathbf{P}^s(N) = \mathbf{P}(N)$
for $k = N-1,...,0$ **do**
  $\hat{\mathbf{x}}^-(k+1) = \mathbf{A}\hat{\mathbf{x}}(k)$
  $\mathbf{P}^-(k+1) = \mathbf{A}\mathbf{P}(k)\mathbf{A}^T + \mathbf{Q}$
  $\mathbf{G}(k) = \mathbf{P}(k)\mathbf{A}^T[\mathbf{P}^-(k+1)]^{-1}$
  $\hat{\mathbf{x}}^s(k) = \hat{\mathbf{x}}(k) + \mathbf{G}(k)[\hat{\mathbf{x}}^s(k+1) - \hat{\mathbf{x}}^-(k+1)]$
  $\mathbf{P}^s(k) = \mathbf{P}(k) + \mathbf{G}(k)[\mathbf{P}^s(k+1) - \mathbf{P}^-(k+1)]\mathbf{G}^T(k)$

Outputs: $\hat{\mathbf{x}}^s(k)$, $\mathbf{P}^s(k)$, $\mathbf{G}(k)$, $k = 1,...,N$

B. Parameters Estimation

We turn to the problem of estimating the parameters η specifying the linear model (17). Adopting the Bayesian perspective, we would like to determine the value of η which maximizes the marginal posterior

$$p(\eta|Y) = \int p(\mathbf{X}, \eta|Y) \ \mathrm{d}\mathbf{X} \tag{33}$$

where $\mathbf{X} := [\mathbf{x}^T(0) \ \cdots \ \mathbf{x}^T(N)]^T$. However, the computation of such a high-dimensional integral is typically avoided by exploiting the decomposition $p(\eta|Y) \propto p(Y|\eta)p(\eta)$ and resorting to an iterative algorithm to compute the likelihood

$$p(Y|\eta) = \int p(\mathbf{X}, Y|\eta) \ \mathrm{d}\mathbf{X}. \tag{34}$$

Here we assume $p(\eta) \propto p(A)p(\alpha)p(\sigma)p(\lambda)$, where $p(\sigma)$ and $p(\lambda)$ are uninformative priors, while $p(\text{vec}(A^T)) \sim \mathcal{N}(\mathbf{0}, \Gamma)$ and $p(\alpha) \sim \mathcal{N}(\mu_\alpha, \Sigma_\alpha)$.
In this work we use the Expectation-Maximization (EM) algorithm to maximize $\ln p(Y|\eta)$, which is equivalent to maximizing $p(Y|\eta)$. The EM method performs such optimization by iteratively maximizing the lower bound

$$\mathscr{F}(q(\mathbf{X}), \eta) = \int q(\mathbf{X}) \ln \frac{p(\mathbf{X}, Y|\eta)}{q(\mathbf{X})} \ \mathrm{d}\mathbf{X} \tag{35}$$

of $\ln p(Y|\eta)$, w.r.t. an arbitrary distribution $q(\mathbf{X})$ and η. In the statistical learning literature $\mathscr{F}(q(\mathbf{X}), \eta)$ is also known as *free-energy*. At the l-th iteration of the algorithm, $\mathscr{F}(q(\mathbf{X}), \eta^{(l)})$ is maximized by $q^{(l+1)}(\mathbf{X}) = p(\mathbf{X}|Y, \eta^{(l)})$. Plugging this into (35), one obtains

$$\mathscr{F}(q^{(l+1)}(\mathbf{X}), \eta) = \int p(\mathbf{X}|Y, \eta^{(l)}) \ln p(\mathbf{X}, Y|\eta) \ \mathrm{d}\mathbf{X} - \int p(\mathbf{X}|Y, \eta^{(l)}) \ln p(\mathbf{X}|Y\eta^{(l)}) \ \mathrm{d}\mathbf{X} \tag{36}$$

which now needs to be maximized w.r.t. η. Noticing that the second term of $\mathscr{F}(q^{(l+1)}(\mathbf{X}),\eta)$ does not depend on η, this coincides with optimizing

$$\mathscr{Q}(\eta,\eta^{(l)}) = \int p(\mathbf{X}|Y,\eta^{(l)}) \ln p(\mathbf{X},Y|\eta)\ \mathrm{d}\mathbf{X} \tag{37}$$

w.r.t. η. Using the Markovian property of system (17), $\mathscr{Q}(\eta,\eta^{(l)})$ can be rewritten as ([21], Ch.12)

$$\begin{aligned}\mathscr{Q}(\eta,\eta^{(l)}) = \sum_{k=1}^{N} \int p(\mathbf{x}(k),\mathbf{x}(k-1)|Y,\eta^{(l)}) \times \\ \times \ln p(\mathbf{x}(k)|\mathbf{x}(k-1),\eta)\mathrm{d}\mathbf{x}(k)\mathrm{d}\mathbf{x}(k-1) \\ + \sum_{k=1}^{N} \int p(\mathbf{x}(k)|Y,\eta^{(l)}) \ln p(y(k)|\mathbf{x}(k),\eta)\mathrm{d}\mathbf{x}(k).\end{aligned} \tag{38}$$

Plugging the smoothing distributions $p(\mathbf{x}(k)|Y,\eta^{(l)}) = \mathscr{N}(\hat{\mathbf{x}}^s(k),\mathbf{P}^s(k))$ (see Eq. (32)) and the pairwise smoothing distributions ([21], Ch.12)

$$\begin{aligned}p(\mathbf{x}(k),\mathbf{x}(k-1)|Y,\eta^{(l)}) = \\ \mathscr{N}\left(\begin{bmatrix}\hat{\mathbf{x}}^s(k)\\ \hat{\mathbf{x}}^s(k-1)\end{bmatrix}, \begin{bmatrix}\mathbf{P}^s(k) & \mathbf{P}^s(k)\mathbf{G}^T(k-1)\\ \mathbf{G}(k-1)\mathbf{P}^s(k) & \mathbf{P}^s(k-1)\end{bmatrix}\right)\end{aligned} \tag{39}$$

into (38) we get

$$\begin{aligned}\mathscr{Q}(\eta,\eta^{(l)}) = &-\frac{N}{2}\ln|2\pi\mathbf{Q}| - \frac{N}{2}\ln|2\pi R| \\ &-\frac{N}{2}\mathrm{tr}\left[\mathbf{Q}^{-1}\left(\Lambda - \Psi\mathbf{A}^T - \mathbf{A}\Psi^T + \mathbf{A}\Upsilon\mathbf{A}^T\right)\right] \\ &-\frac{N}{2}\mathrm{tr}\left[R^{-1}\left(\Delta - \Xi\mathbf{C}^T - \mathbf{C}\Xi^T + \mathbf{C}\Lambda\mathbf{C}^T\right)\right]\end{aligned} \tag{40}$$

where

$$\begin{aligned}\Lambda &= \frac{1}{N}\sum_{k=1}^{N} \mathbf{P}^s(k) + \hat{\mathbf{x}}^s(k)\left[\hat{\mathbf{x}}^s(k)\right]^T \\ \Upsilon &= \frac{1}{N}\sum_{k=1}^{N} \mathbf{P}^s(k-1) + \hat{\mathbf{x}}^s(k-1)\left[\hat{\mathbf{x}}^s(k-1)\right]^T \\ \Xi &= \frac{1}{N}\sum_{k=1}^{N} y(k)\left[\hat{\mathbf{x}}^s(k)\right]^T \\ \Psi &= \frac{1}{N}\sum_{k=1}^{N} \mathbf{P}^s(k)\mathbf{G}(k-1) + \hat{\mathbf{x}}^s(k)\left[\hat{\mathbf{x}}^s(k-1)\right]^T \\ \Delta &= \frac{1}{N}\sum_{k=1}^{N} y(k)y^T(k).\end{aligned}$$

The complete routine we adopt to estimate the parameters η through the EM method is summarized in Algorithm 3. The procedure is initialized with $\eta^{(0)} = \{A^{(0)},\sigma^{(0)},\alpha^{(0)},\lambda^{(0)}\}$:

$$\begin{aligned}A^{(0)} &= -I_n, \qquad \sigma^{(0)} = 10^{-2} \\ \alpha^{(0)} &= \mu_\alpha, \qquad \lambda^{(0)} = \sqrt{\frac{\mathrm{tr}[(Y-\bar{Y})(Y-\bar{Y})^\top]}{10Nn}}\end{aligned} \tag{41}$$

where $\bar{Y} = \frac{1}{Nn}\sum_{i=1}^{Nn} Y_i$.

Notice that, at Step 4 of Algorithm 3, the search of η is restricted to the set $\Omega = \{\mathscr{A},\ \mathbb{R}_+,\ \mathbb{R}^{p+1},\ \mathbb{R}_+\}$. Since this optimization may be computationally intensive, we have also tested a simpler method, where the uncertainty of the estimates returned by the smoother (encoded in the matrix $\mathbf{P}^s(k)$) is ignored. Briefly, this alternative routine iteratively estimates A and σ through Algorithm 1, α and λ through Eqs. (30), (31), replacing the hidden neuronal states with the smoother estimates $\hat{\mathbf{x}}^s(k)$. The results are inferior to those of Algorithm 3 and thus not reported.

Algorithm 3 Estimation of parameters η through EM

Inputs: $y(k),\ k=1,..,N$, H, μ_α, Σ_α
Initialization: Set $\eta^{(0)}$ as in Eq. (41), set $l=0$
1: **repeat**
2: Apply Algorithm 2 to get $\hat{\mathbf{x}}^s(k)$, $\mathbf{P}^s(k)$, $\mathbf{G}(k)$, $k=1,...,N$
3: Compute $\mathscr{Q}(\eta,\eta^{(l)})$ using Eq. (40)
4: $\eta^{(l+1)} = \arg\max_{\eta\in\Omega}\Big\{\mathscr{Q}(\eta,\eta^{(l)}) - \frac{1}{2}\ln|2\pi\Gamma^{(l)}| - \frac{1}{2}\ln|2\pi\Sigma_\alpha| - \frac{1}{2}\mathrm{vec}(A^T)^T[\Gamma^{(l)}]^{-1}\mathrm{vec}(A^T) - \frac{1}{2}(\alpha-\mu_\alpha)\Sigma_\alpha^{-1}(\alpha-\mu_\alpha)\Big\}$
5: $Q = \sigma^{(l+1)^2}\int_0^{T_R} e^{A^{(l+1)}\tau}e^{A^{(l+1)^T}\tau}d\tau$
6: $\gamma_i^{(l+1)} = (a_i^{(l+1)})^2 + \gamma_i^{(l)} - (\gamma_i^{(l)})^2\phi_i^T(\Phi\Gamma^{(l)}\Phi^T + Q\otimes I_N)^{-1}\phi_i$, $i=1,...,n^2$
7: $l = l+1$
8: **until** $\|A^{(l)} - A^{(l-1)}\|_F / \|A^{(l)}\|_F$ is sufficiently small
Outputs: $\eta^{(l)}$

VI. Experimental Results

The methods illustrated in Secs. IV and V are experimentally evaluated using simulated rs-fMRI data. In particular: (a) we analyze the effectiveness of the reweighted ℓ_1 technique detailed in Sec. IV-A in detecting the correct brain network structure and in estimating the strength of the neuronal connections; (b) we evaluate the performance of Algorithm 3 which only uses the resting state BOLD signals, when the true network structure has to be inferred from the fMRI data. In the latter setup, the proposed methods are compared with the Variational Bayes routines implemented in the SPM12 toolbox (`http://www.fil.ion.ucl.ac.uk/spm/`).

A. Experimental Setup

The rs fMRI data used in these experiments are generated through the SPM12 routines `spm_int_J`, `spm_fx_fmri` and `spm_gx_fmri`. Specifically, we consider a brain network with $n=7$ neuronal nodes, whose activity is modulated according to Eq. (2), where $w(t)$ is a Gaussian white noise process with variance $\sigma^2 = 0.01$ and

$$A = \begin{bmatrix}-0.5 & 0 & 0 & 0 & -0.2 & 0 & 0\\ 0 & -0.5 & 0 & -0.45 & -0.3 & 0 & 0\\ 0 & 0 & -0.5 & 0.8 & 0 & 0 & 0\\ 0 & 0.6 & 0 & -0.5 & -0.1 & 0.6 & 0\\ 0.3 & 0 & -0.55 & 0 & -0.5 & 0.2 & 0\\ 0 & 0 & 0 & 0 & 0.3 & -0.5 & 0.45\\ 0.15 & 0 & 0.2 & 0 & 0 & 0 & -0.5\end{bmatrix} \tag{42}$$

$N=600$ data are generated with sampling time $T_R = 2$s. No observation noise is added to the BOLD signal generated

	$ERR(\hat{A})$	$RMSE(\hat{A})$
Algorithm 1	3 (2.9)	0.05 (0.11)
Linear reweighted ℓ_1 algorithm	10 (1.4)	0.18 (0.006)

TABLE I: Performance of Algorithm 1 and classical reweighted ℓ_1 method over 50 Monte-Carlo runs (using measured neuronal activity $x(k),\ k = 1,..,600$. Median and std (between brackets) of metrics (43) and (44) are reported.

by the routine `spm_gx_fmri`. It is worth stressing that the self-connections are a-priori fixed to -0.5 to prevent instability issues.

B. Performance Evaluation

Let $\hat{A}$ be an estimate of the effective connectivity matrix: we define the Root-Mean-Squared-Error (RMSE)

$$RMSE(\hat{A}) = \frac{\|\underline{A} - \underline{\hat{A}}\|_F}{\sqrt{n(n-1)}} \tag{43}$$

where the notation $\underline{A}$ denotes the matrix A with its diagonal set to 0. Indeed, the diagonal elements of A (i.e. the self-connections) do not give any information on the effective connectivity.
A second performance index measures the number of errors committed by a certain method in the recovery of the true sparsity pattern of matrix (42). Namely,

$$ERR(\hat{A}) = \|SP(A) - SP(\hat{A})\|_F^2 \tag{44}$$

where $SP(\cdot) : \mathbb{R}^{n\times m} \to \mathbb{R}^{n\times m}$ and

$$[SP(A)]_{ij} := \begin{cases} 1 & \text{if } [A]_{ij} \neq 0 \\ 0 & \text{if } [A]_{ij} = 0 \end{cases} . \tag{45}$$

C. Estimation of effective connectivity using measured neuronal activity

In this section we assess the validity of the reweighted ℓ_1 method illustrated in Sec. IV-A. To this purpose we pass as inputs to Algorithm 1 the neuronal time-series $x(k)$, $k = 1,...,600$, generated by the SPM12 routine `spm_int_J` when provided with `spm_fx_fmri` and the matrix (42). In addition, we compare Algorithm 1 with the linear reweighted ℓ_1 based on approximation (23).
Table I compares the results achieved by Algorithm 1 and by its linear version on the data generated in 50 Monte-Carlo runs, using different realizations of the noise $w(t)$: the medians (and the standard deviations) of indexes (43) and (44) are reported. Before evaluating these indexes, a tresholding is applied to the estimated matrices $\hat{A}$: namely the entries with absolute value less than 0.1 are set to zero. The results reported in Table I show that Algorithm 1 benefits from using the true generative model at Step 4, in place of the linear approximation (23).

D. Estimation of effective connectivity using rs fMRI data

We now evaluate the performance of the proposed Algorithm 3 in the recovery of matrix (42) when the BOLD data $y(k),\ k = 1,...,600$, generated under the setup of Sec. VI-A, are available. In particular, we conduct a Monte-Carlo study over 20 different realizations of the noise $w(k)$. The performance of Algorithm 3 is compared with the one returned by the SPM12 routines `spm_dcm_fmri_csd` and `spm_dcm_estimate`, both provided with null exogenous inputs. These functions implement two of the methods mentioned in Sec. II-B: specifically, `spm_dcm_fmri_csd` runs the spectral DCM (**sDCM**) [6], while `spm_dcm_estimate` is invoked activating the option `DCM.options.stochastic` in order to implement the **DEM** procedure [8].
To guarantee a fair comparison between our methods and the SPM12 routines, the latter should perform model selection among all the possible sparsity patterns of matrix A. However, this would require the estimation of a combinatorial number of matrices, making this type of approach computationally infeasible. Attempting to still conduct a fair comparison, we pass 21 candidate sparsity patterns m_i, $i = 1,...,21$, to the SPM12 routines, including the true one (denoted by **(a)**). These are illustrated in Figure 2. For each of them, Table II reports the value of index (44), as well as the median and the standard deviation of index (43) corresponding to the matrices estimated by the routines `spm_dcm_fmri_csd` and `spm_dcm_estimate`. In addition, the columns denoted with *#Chosen* report the number of times each model is chosen according to the free-energy metric. The most frequently selected structures are very different from the true one (as indicated by the value of the metrics $ERR(\hat{A})$). According to our experimental experience, the SPM routines tend to select structures with a low degree of sparsity.
Table III contains the median and the standard deviation of indexes (43) and (44) computed for the matrices $\hat{A}$ estimated by the proposed Algorithm 3. The performance of our approach is superior to the one achieved by the SPM12 routines, both in terms of $RMSE(\hat{A})$ and $ERR(\hat{A})$. In addition, the sparsity patterns detected by Algorithm 3 differ at most by 7 entries from the true one. It is worth noting that $RMSE(\hat{A})$ returned by Algorithm 3 is comparable to that achieved by spectral DCM when the latter is provided with the true sparsity pattern.
On the conducted study, Alg. 3 takes an average computational time equal to 558s, while `spm_dcm_fmri_csd` and `spm_dcm_estimate` respectively take on average 721s and 4044s to process all the 21 patterns.

VII. Conclusions and Future Work

We have proposed a new approach for the estimation of the effective connectivity in brain networks using resting state fMRI data. Existing methods tackle the problem by formulating a parametric generative model (DCM) of the observed data and by inverting it using Bayesian approaches. However, due to the nonlinear and stochastic nature of the

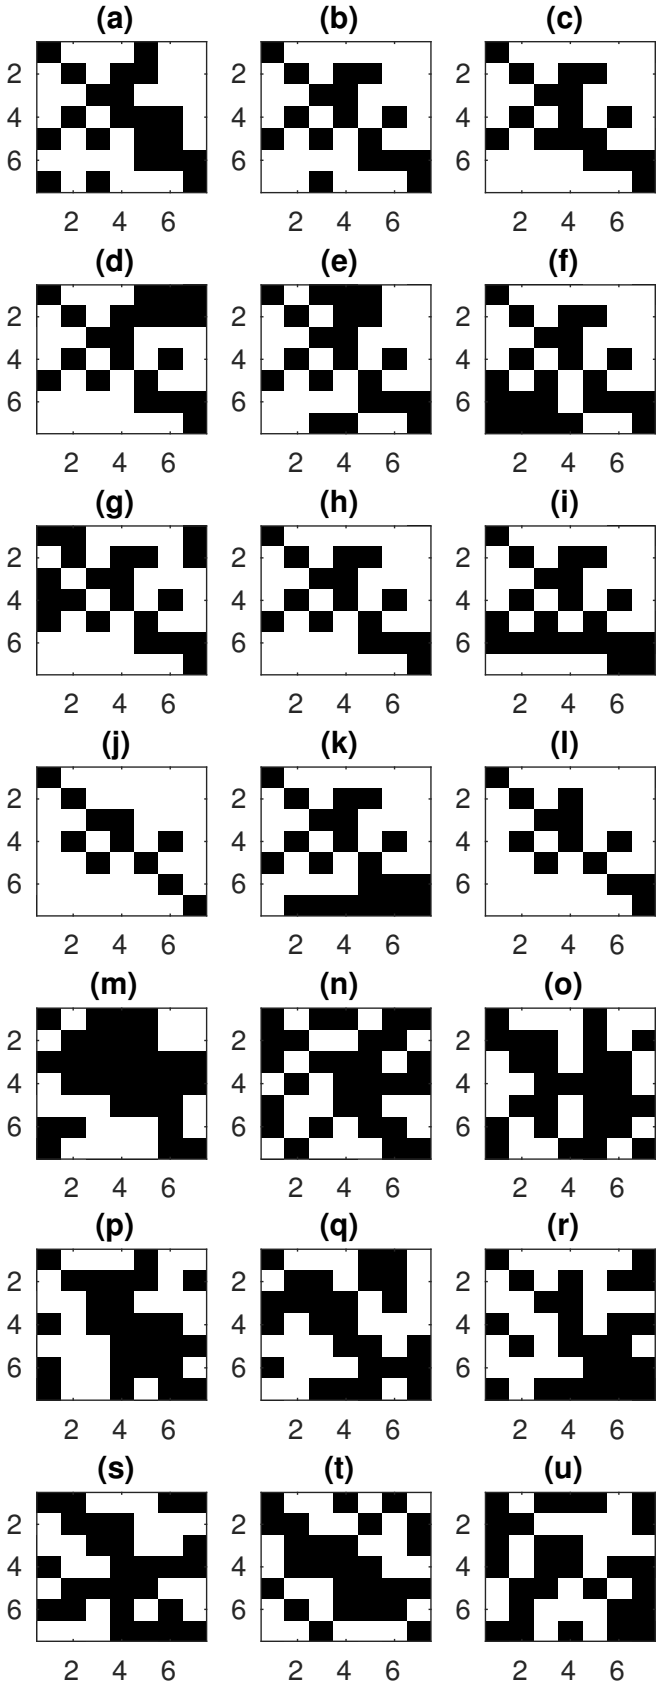

Fig. 2: Candidate sparsity patterns m_i which are provided to the SPM12 routines. Black cells denote non-zero elements. **(a)** is the true one.

		spectral DCM		**DEM**	
	$ERR(\hat{A})$	*#Chosen*	$RMSE(\hat{A})$	*#Chosen*	$RMSE(\hat{A})$
(a)	0	2	0.12 (0.05)	0	0.21 (0.004)
(b)	4	0	0.25 (0.02)	0	0.21 (0.005)
(c)	6	0	0.23 (0.03)	0	0.21 (0.003)
(d)	8	0	0.18 (0.02)	0	0.21 (0.004)
(e)	6	1	0.14 (0.05)	0	0.21 (0.005)
(f)	8	0	0.22 (0.04)	0	0.21 (0.005)
(g)	10	0	0.28 (0.09)	0	0.22 (0.007)
(h)	5	0	0.22 (0.02)	0	0.20 (0.003)
(i)	10	0	0.24 (0.07)	0	0.21 (0.004)
(j)	10	0	0.17 (0.004)	0	0.22 (0.002)
(k)	8	0	0.27 (0.03)	0	0.21 (0.005)
(l)	8	0	0.18 (0.01)	0	0.22 (0.002)
(m)	19	1	0.33 (0.04)	3	0.23 (0.003)
(n)	**22**	**14**	**0.31 (0.06)**	2	0.23 (0.005)
(o)	**19**	0	0.36 (0.04)	**14**	**0.23 (0.003)**
(p)	15	0	0.33 (0.04)	0	0.22 (0.002)
(q)	21	2	0.32 (0.04)	0	0.23 (0.003)
(r)	16	0	0.28 (0.02)	0	0.22 (0.002)
(s)	24	0	0.31 (0.03)	0	0.23 (0.004)
(t)	17	0	0.28 (0.03)	1	0.23 (0.002)
(u)	23	0	0.27 (0.02)	0	0.22 (0.002)

TABLE II: Performance of the SPM12 routines `spm_dcm_fmri_csd` (**sDCM**) and `spm_dcm_estimate` (**DEM**) over 20 Monte-Carlo runs, when they are given the 21 sparsity patterns of Fig. 2. Median and std (within brackets) of index (43) are reported, together with the no. of times (*#Chosen*) each sparsity pattern is chosen according to the free-energy metric. Bold font highlights the most frequently selected models.

	$ERR(\hat{A})$	$RMSE(\hat{A})$
Algorithm 3	4 (0.74)	0.13 (0.008)

TABLE III: Performance of Algorithm 3 over 20 Monte-Carlo runs. Median and std (within brackets) of the metrics (43) and (44) are reported.

standard DCMs, the current techniques are highly computationally intensive. In addition, when no a-priori knowledge is available on the structure of the underlying network connecting the monitored brain regions, these approaches require to compare several candidate interaction patterns in order to choose the most plausible one.

The technique we have developed is based on a linear stochastic generative model (DCM), which we have built by linearly approximating the Balloon model describing the hemodynamic response. Exploiting the linearity of the derived DCM we have applied classical smoothing techniques to estimate both the hidden neuronal activity and the model parameters starting from the observed BOLD signal. A key feature of our algorithms is the ability to automatically detect the structure of the underlying brain network, thanks to the use of a sparsity inducing method, thus avoiding a combinatorial search over candidate structures.

Future work will include an extended experimental study of the proposed method and, in particular, its application to real fMRI data.

REFERENCES

[1] K. J. Friston, "Functional and effective connectivity: a review," *Brain connectivity*, vol. 1, no. 1, pp. 13–36, 2011.

[2] K. J. Friston, L. Harrison, and W. Penny, "Dynamic causal modelling," *Neuroimage*, vol. 19, no. 4, pp. 1273–1302, 2003.

[3] A. C. Marreiros, S. J. Kiebel, and K. J. Friston, "Dynamic causal modelling for fmri: a two-state model," *Neuroimage*, vol. 39, no. 1, pp. 269–278, 2008.

[4] K. E. Stephan, L. Kasper, L. M. Harrison, J. Daunizeau, H. E. den Ouden, M. Breakspear, and K. J. Friston, "Nonlinear dynamic causal models for fmri," *Neuroimage*, vol. 42, no. 2, pp. 649–662, 2008.

[5] B. Li, J. Daunizeau, K. E. Stephan, W. Penny, D. Hu, and K. Friston, "Generalised filtering and stochastic dcm for fmri," *Neuroimage*, vol. 58, no. 2, pp. 442–457, 2011.

[6] K. J. Friston, J. Kahan, B. Biswal, and A. Razi, "A dcm for resting state fmri," *Neuroimage*, vol. 94, pp. 396–407, 2014.

[7] K. Friston, J. Mattout, N. Trujillo-Barreto, J. Ashburner, and W. Penny, "Variational free energy and the laplace approximation," *Neuroimage*, vol. 34, no. 1, pp. 220–234, 2007.

[8] K. J. Friston, N. Trujillo-Barreto, and J. Daunizeau, "Dem: a variational treatment of dynamic systems," *Neuroimage*, vol. 41, no. 3, pp. 849–885, 2008.

[9] J. Daunizeau, K. J. Friston, and S. J. Kiebel, "Variational bayesian identification and prediction of stochastic nonlinear dynamic causal models," *Physica D: Nonlinear Phenomena*, vol. 238, no. 21, pp. 2089–2118, 2009.

[10] K. Friston and W. Penny, "Post hoc bayesian model selection," *Neuroimage*, vol. 56, no. 4, pp. 2089–2099, 2011.

[11] R. B. Buxton, E. C. Wong, and L. R. Frank, "Dynamics of blood flow and oxygenation changes during brain activation: the balloon model," *Magnetic resonance in medicine*, vol. 39, no. 6, pp. 855–864, 1998.

[12] K. J. Friston, A. Mechelli, R. Turner, and C. J. Price, "Nonlinear responses in fmri: the balloon model, volterra kernels, and other hemodynamics," *NeuroImage*, vol. 12, no. 4, pp. 466–477, 2000.

[13] K. E. Stephan, N. Weiskopf, P. M. Drysdale, P. A. Robinson, and K. J. Friston, "Comparing hemodynamic models with dcm," *Neuroimage*, vol. 38, no. 3, pp. 387–401, 2007.

[14] K. Friston, K. Stephan, B. Li, and J. Daunizeau, "Generalised filtering," *Mathematical Problems in Engineering*, vol. 2010, 2010.

[15] A. Razi, J. Kahan, G. Rees, and K. J. Friston, "Construct validation of a dcm for resting state fmri," *Neuroimage*, vol. 106, pp. 1–14, 2015.

[16] S. Frässle, E. I. Lomakina, A. Razi, K. J. Friston, J. M. Buhmann, and K. E. Stephan, "Regression dcm for fmri," *NeuroImage*, 2017.

[17] D. Wipf and S. Nagarajan, "Iterative reweighted ℓ_1 and ℓ_2 methods for finding sparse solutions," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 317–329, 2010.

[18] D. P. Wipf and S. S. Nagarajan, "A new view of automatic relevance determination," in *Advances in neural information processing systems*, 2008, pp. 1625–1632.

[19] M. E. Tipping, "Sparse bayesian learning and the relevance vector machine," *Journal of machine learning research*, vol. 1, no. Jun, pp. 211–244, 2001.

[20] H. E. Rauch, C. Striebel, and F. Tung, "Maximum likelihood estimates of linear dynamic systems," *AIAA journal*, vol. 3, no. 8, pp. 1445–1450, 1965.

[21] S. Särkkä, *Bayesian filtering and smoothing*. Cambridge University Press, 2013, vol. 3.